arXiv:1609.09728v1 [cs.FL] 30 Sep 2016

On the Complexity of Bounded Context Switching

Peter Chini, Jonathan Kolberg, Roland Meyer, Prakash Saivasan

TU Kaiserslautern, {chini, kolberg, meyer, saivasan}@cs.uni-kl.de

Abstract

Bounded context switching (BCS) is an under-approximate method for finding violations to safety properties in shared memory concurrent programs. Technically, BCS is a reachability problem that is known to be NP-complete if the context bound is part of the input. Our contribution is a parameterized analysis of BCS. The main finding is a new algorithm that solves BCS when parameterized by the number of context switches and the size of the memory. Notably, there are no restrictions on the number, the size, nor the behavior of the threads. Moreover, the algorithm is generic in that it applies to various classes of programs. The runtime is $\mathcal{O}^*(m^{cs} \cdot 4^{cs})$ where m is the size of the memory. We show an $\mathcal{O}^*(2^{cs})$ lower bound, assuming the exponential time hypothesis. The idea behind the algorithm is to iterate over the memory sequences: sequences of states in the memory that are seen upon context switches. We also study a generalization of the parameterization that only limits the size of the memory each thread may access.

Keywords: Shared memory concurrency, safety verification, formal languages, fixed parameter tractability, exponential time hypothesis.

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Concurrent programs are present at different levels of the software stack, from high-level applications to low-level software implementing system services on multicores. A concurrent program consists of a number of threads that communicate via a shared memory holding values from some data domain. Concurrent programs are hard to get right, which makes it important to develop verification methods for checking their correctness. Safety properties require the computations of a program of interest to stay within a given set of safe states. In fact, most of the requirements imposed on concurrent programs are safety properties. The design of automatic verification methods for safety properties of concurrent programs remains a challenge. In practice, the interactions between the threads can lead to subtle safety violations that are hard to predict and to reproduce. In terms of complexity[1], when the threads are recursive the safety verification problem is undecidable, even for two threads communicating through one bit of shared memory [2]. In the finite state setting, the problem is PSPACE-complete [21].

To reduce the complexity, a trend in verification is to give up completeness and consider approximate analysis algorithms. For safety properties, under-approximate analyses have received considerable attention. The idea is bug hunting: under-approximations explore a restricted set of computations in order to find an unsafe state. Hence, they are sound but only complete if all computations of the program happen to match the restricted set. The quality of an under-approximation is determined by two aspects. The restricted behavior should be (i) likely to contain a bug if one exists and (ii) algorithmically easy to explore.

[1] For a complexity analysis of verification problems, it is common to assume both the number of threads and the data domain to be finite.

licensed under Creative Commons License CC-BY

A popular restriction that was introduced in [27] is *bounded context switching (BCS)*. A context switch occurs if a thread gives away the processor for another thread to be scheduled. Bounding the context switches thus limits the communication capabilities, the number of times the threads can communicate via the shared memory. Studies have shown that many violations to safety properties occur within few context switches [26]. Moreover, the complexity of the analysis drops from undecidable to being NP-complete [27].

In this work, we conduct a parameterized complexity analysis of bounded context switching. The hope is that the problem, despite being NP-hard, can still be solved efficiently with a non-polynomial blow-up only in controllable parameters. More formally, we seek for *fixed parameter tractability (FPT)* results: The complexity of the problem should be $f(k) \cdot p(n)$ where $f(k)$ is a function only in the controllable parameter k and $p(n)$ is a polynomial in the size of the input. The first finding is that BCS when parameterized by the context switches is W[1]-complete. W[1]-complete problems are believed to be strictly harder than FPT, which motivates the search for further parameterizations.

Our main finding is an FPT algorithm for BCS when parameterized by the number of context switches *and* the size of the main memory. Notably, there are no restrictions on the number, the size, nor the behavior of the threads. Moreover, the algorithm carries over to more general programming models, in particular the recursive case. The runtime is $\mathcal{O}^*(m^{cs} \cdot 4^{cs})$, where m is the size of the memory. We show a $\mathcal{O}^*(2^{cs})$ lower bound, assuming the exponential time hypothesis.

The algorithm proceeds in two phases. First, it decomposes the problem into $2^{cs \cdot \log m}$ subproblems based on the following observation. A concurrent computation that reaches an unsafe state yields a memory sequence, a sequence of states in the shared memory that are seen upon context switches. Hence, we can iterate over all memory sequences and check whether one of them is the result of a concurrent computation. This check is conducted in the second phase. We give a graph construction that, conceptually, covers a given memory sequence with executions from the different threads.

The next step is to relax the parameterization and only limit the memory locally, in terms of the amount that each thread may access. Here, we again face W[1]-hardness. If we make the number of threads a parameter, we recover FPT by a bounded search tree. Indeed, FPT follows from the previous construction but in this setting the bounded search tree yields a better complexity.

Related work. Verification of concurrent systems has received considerable attention. The parameterized verification of asynchronous shared memory systems has been studied in [12, 13, 18, 19]. The model considered there is that of a distinguished leader and unboundedly many contributors. The authors study the question whether safety properties hold no matter how many contributors are present. The problem is shown to be coNP-complete when only finite state systems are involved and PSPACE-complete when recursive threads are involved [13]. In our work, we consider verification of safety properties for a fixed number of threads, which is known to be undecidable in their setting.

In [2], the safety verification of shared memory systems under a *bounded-stage restriction* with fixed number of threads was analyzed. It was shown that under such a restriction, when only finite processes are involved, the problem is NP-complete, NEXP-time when counter-systems are involved and it is undecidable when recursive processes are involved. In [4], the safety verification of shared memory systems with a fixed number of threads was analyzed when the information flow between the threads was restricted to be acyclic. Our work involves analyzing the problem under a *bounded-context restriction* to provide a *fixed parameter tractable* solution.

The idea of bounded context was introduced in [27] and reachability on concurrent systems under this restriction was shown to be NP-complete. It was shown in [26] that in most practical cases, context bounding is an effective way to capture bugs. In [22, 24], it was shown that context bounded reachability problems on concurrent systems can be reduced to reachability on sequential systems. In [3] concurrent systems with the ability to dynamically create threads were considered and bounded context reachability under this setting was shown to be EXPSPACE-complete. In [23], the parametrized verification problem of concurrent systems was solved by composing the interfaces of k-bounded context runs for increasing values of k.

Besides introducing the notion of conflict serializability under TSO, the authors of [1] gave a fixed parameter tractable algorithm for checking serializability. In [9], various model checking problems for synchronized executions on parallel components were considered and proven to be intractable. Parameterized complexity analyses for two problems on automata were given in [15]. The authors give classifications in terms of higher intractability classes and prove simple tractable algorithms to be optimal. Also in [28], a complete parameterized complexity analysis of the intersection non-emptiness problem is shown. Our work provides a parameterized complexity analysis for bounded context switching.

Lower bounds, based on the *(strong) exponential time hypothesis*, on the runtime of algorithms related to automata were already considered in [16]. The authors gave lower bounds for the intersection non-emptiness problem, and for universality/equivalence. Our work adds a lower bound for bounded context switching.

2 Bounded Context Switching and Fixed Parameter Tractability

We define the bounded context switching problem of interest and recall the basics on fixed parameter tractability following [11, 17]. Furthermore, we present a first result that displays the complexity of a variant of bounded context switching.

2.1 Bounded Context Switching

We consider the safety verification problem for shared memory concurrent programs. Such a program consists of different threads and a shared memory that is used for communication activities between the threads. To obtain precise complexity results, it is common to assume both the number of threads and the data domain to be finite. Safety properties partition the states of a program into *unsafe* and *safe* states. Hence, checking safety amounts to checking whether no *unsafe* state is reachable. In the following, we develop a language-theoretic formulation of the reachability problem that will form the basis of our study.

We model the shared memory as a (non-deterministic) finite automaton of the form $M = (Q, \Sigma, \delta_M, q_0, q_f)$. The states Q correspond to the data domain, the set of values that the memory can be in. The initial state $q_0 \in Q$ is the value that the computation starts from. The final state $q_f \in Q$ reflects the reachability problem. The alphabet Σ models the set of operations. Operations have the effect of changing the memory valuation, formalized by the labeled transition relation $\delta_M \subseteq Q \times \Sigma \times Q$ (we use $q \xrightarrow{a} q'$ to denote $(q, a, q') \in \delta_M$). We generalize the transition relation to words $u \in \Sigma^*$. The set of valid sequences of operations that lead from a state q to another state q' is the language $L(A(q, q')) := \{u \in \Sigma^* \mid q' \in \delta_M(q, u)\}$. The language of A is $L(A) := L(A(q_0, q_f))$.

We also model the threads operating on the shared memory M as finite automata $A = (P, \Sigma \times \{A\}, \delta_A, p_0, p_f)$. Note that they use the alphabet Σ of the shared memory, indexed by the name of the thread. The index will play a role when we define the notion

of context switches below. The automaton A is nothing but the control flow graph of the thread. Its language is the sequences of operations that the thread may potentially execute to reach the final state. Not all these sequences will be feasible, the reason being that the thread language does not take into account the effect of the operations on the shared memory. Indeed, the thread may issue a command $write(x, 1)$ followed by $read(x, 0)$, which the automaton for the shared memory will reject. The computations of A that are actually feasible on the shared memory are given by the intersection $L(M) \cap L(A)$. Here, we silently assume the intersection to project away the second component of the thread alphabet.

In a concurrent program, we have multiple threads A_1 to A_t that mutually influence each other by accessing the same memory M. We mimic this influence by interleaving the thread languages, formalized with the shuffle operator ⧢. For the definition, consider languages $L_1 \subseteq \Sigma_1^*$ and $L_2 \subseteq \Sigma_2^*$ over disjoint alphabets $\Sigma_1 \cap \Sigma_2 = \emptyset$. The shuffle of the languages $L_1 \text{⧢} L_2 := \{u \in (\Sigma_1 \cup \Sigma_2)^* \mid u \downarrow \Sigma_i \in L_i \cup \{\varepsilon\}, i = 1, 2\}$ contains all words over the union of the alphabets where the corresponding projections $(- \downarrow -)$ belong to the operand languages.

We are now prepared to define a *shared memory concurrent program (SMCP)* as a tuple $S = (\Sigma, M, (A_i)_{i\in[1..t]})$. The language of the program is $L(S) := L(M) \cap (\, \text{⧢}_{i\in[1..t]} L(A_i)\,)$, and the safety verification problem induced by the program is to decide whether $L(S)$ is non-empty. This problem is algorithmically hard due to the complex interactions among the threads.

We formulate the notion of context switching in our setting. Every word in the shuffle of the thread languages, $u \in \text{⧢}_{i\in[1..t]} L(A_i)$, has a unique decomposition into maximal infixes that are generated by the same thread. Formally, $u = u_1 \dots u_{cs+1}$ so that there is a function $\varphi : [1..cs+1] \to [1..t]$ satisfying $u_i \in (\Sigma \times \{A_{\varphi(i)}\})^+$ and $\varphi(i) \neq \varphi(i+1)$ for all $i \in [1..cs]$. We refer to the u_i as contexts and to the thread changes between u_i to u_{i+1} as *context switches*. So $u = u_1 \dots u_{cs+1}$ has $cs+1$ contexts and cs context switches. Let $\text{Context}(\Sigma, t, cs)$ denote the set of all words (over Σ with t threads) that have cs-many context switches.

The *bounded context switching* under-approximation limits the safety verification task to the set of all computations in S that have cs context switches. Formally, we define the language of cs-bounded words to be $L(S, cs) := L(S) \cap \text{Context}(\Sigma, t, cs)$ and define the problem as follows.

Bounded Context Switching (BCS)	
Input:	An SMCP $S = (\Sigma, M, (A_i)_{i\in[1..t]})$ and a bound $cs \in \mathbb{N}$.
Question:	Is $L(S, cs) \neq \emptyset$?

2.2 Fixed Parameter Tractability

The algorithmic result that we start from is the NP-completeness of BCS due to [14] (which holds even in the case of an unary alphabet). Our goal is to understand the dependency of the analysis on the parameter cs. To this end, we consider parametrizations of BCS and employ a parameterized complexity analysis. The idea is to show that BCS, despite its NP-hardness, is still solvable efficiently. This means that the runtime of an algorithm splits into a polynomial and a non-polynomial part and only cs, and other chosen parameters, are allowed to occur in the non-polynomial part.

Formally, a parameterized problem is called *fixed parameter tractable (FPT)* if it can be solved by a deterministic algorithm that runs in time $f(k) \cdot n^c$, where n is the size of the input, c is a constant and f is a computable function that only depends on the parameter k. It is common to denote the runtime by $\mathcal{O}^*(f(k))$ to suppress the polynomial part.

While many parametrizations of NP-hard problems were proven to be fixed parameter tractable, there are problems that are unlikely to be FPT. A famous example is the problem of finding a *clique* of size k: despite comprehensive effort, no $f(k) \cdot n^c$-time algorithm was found for the problem. In fact, it is unlikely that such an algorithm exists since the *clique*-problem is complete for W[1], a class of *intractable* problems.

To establish a theory of intractable problems, it is essential to define *reductions*. Given two parameterized problems L and L', we say that L is *reducible* to L' via a *parameterized reduction* if there is an algorithm that transforms an instance (x, k) of L to an instance (x', k') of L' in time $g(k) \cdot n^d$, where g is any computable function and d is a constant, so that $(x, k) \in L$ if and only if $(x', k') \in L'$. Note that the parameter k' is a function only depending on k. If L is reducible to L' via a parameterized reduction, we denote this by $L \leq^{fpt} L'$. The basic class of intractable problems is denoted by W[1]. We may think of it as the class of all problems that are reducible (via a parameterized reduction) to the problem of finding a satisfying assignment with exactly k literals set to true for a given 3-CNF formula. We have that FPT $\subseteq$ W[1] and it is believed that the inclusion is strict.

The runtime of a fixed parameter tractable algorithm is dominated by the function f. A goal of recent algorithm design is to develop new approaches to improve on f. But there is also a contrary approach based on the *Exponential Time Hypothesis (ETH)*. This is the assumption that there is no $2^{o(n)}$-time algorithm for solving n-variable 3-SAT [20]. For particular problems one can prove that f cannot drop below a certain lower bound unless ETH fails. In the last years, the approach has proven to be useful and various lower bounds for different problems were discovered [7, 8, 25].

Our parameterized complexity analysis of BCS begins with an intractability result. We denote by BCS(cs) and BCS(cs, t) two versions of the problem that are identical to BCS except that we introduce parameters. BCS(cs) has as a parameter the number of context switches while BCS(cs, t) is parameterized by both, the context switches and the number of threads. The next theorem shows that both parametrizations are intractable. The proof can be found in Appendix A.

▶ **Theorem 1.** BCS(cs) *and* BCS(cs, t) *are both* W[1]*-complete.*

The theorem shows that we have to consider alternative parametrizations and variants of BCS in order to obtain a fixed parameter tractable algorithm. This will be our guideline for the next sections.

3 Bounded Memory

We present a complexity analysis of BCS that is parameterized by the size of the memory and hence the amount of information the threads may exchange in one communication. We chose the size for two reasons. First, bounding the size is a liberal restriction in the sense that it neither limits the number of threads, nor their size, nor their protocol of interaction. Moreover, threads often communicate by signaling, which does not need much information. Therefore, one may indeed hope to find bugs in this restricted set of computations (if the program is incorrect). Second, as we shall see, parameterizing by the size of the memory allows us to devise an efficient algorithm for checking non-emptiness. Interestingly, the algorithm carries over to more general programming models, in particular the recursive case. We also provide a lower bound based on the exponential time hypothesis. Together, the parameterization meets the requirements on a useful under-approximation that we discussed above.

3.1 Fixed Parameter Tractability

We define the following parameterized variant of BCS, denoted by BCS$(cs, |Q|)$. Let the input be an SMCP $S = (\Sigma, M, (A_i)_{i\in[1..t]})$ and an integer $cs \in \mathbb{N}$. The parameters are the number of context switches cs and the number of states $|Q|$ in the given memory automaton M. Our contribution is an FPT algorithm (in these parameters). In the statement of Theorem 2, we give the precise influence of both parameters rather than taking their maximum. Moreover, we use $|P|$ to denote the maximal number of states in the thread automata $A_1, \ldots, A_t$.

▶ **Theorem 2.** BCS$(cs, |Q|)$ *can be solved in time* $\mathcal{O}(|Q|^{cs} \cdot 4^{cs} \cdot t \cdot cs \cdot |P| + t \cdot |Q|^3 \cdot |P|^3)$.

The algorithm proceeds in two phases: it decomposes the problem into exponentially many simpler problems and solves each of them. To decompose the problem, we define *memory sequences*, sequences of states (from Q^*) that the memory automaton is in upon a context switch. The reasoning is as follows. If the program is unsafe, which means $L(S) \neq \emptyset$, there is a word in the shuffle of the thread languages that is also accepted by the memory automaton. Hence, the word induces (at least one) memory sequence. Formally, $u \in L(S)$ *induces a memory sequence* $w \in Q^*$ if there is an accepting run of M on u where, upon every context switch, the state of M is the one given by w. We use $M(S) \subseteq Q^*$ for the language of all induced memory sequences. The definition yields:

▶ **Lemma 3.** $L(S) \neq \emptyset$ *if and only if* $M(S) \neq \emptyset$.

With this equivalence, we can turn around the reasoning and understand an induced memory sequence as a witness of non-emptiness for $L(S)$. Since the number of context switches is given by the parameter cs, we can iterate over all memory sequences in Q^{cs} and check each of them for being a witness, an *induced* memory sequence. Hence, $|Q|^{cs}$ appears in the complexity.

The simpler problem resulting from the decomposition is the following.

Witness (WIT)	
Input:	An SMCP $S = (\Sigma, M, (A_i)_{i\in[1..t]})$ and a memory sequence $w \in Q^{cs}$.
Question:	Is $w \in M(S)$?

We develop an algorithm that solves WIT in time $\mathcal{O}(4^{cs} \cdot t \cdot cs \cdot |P| + t \cdot |Q|^3 \cdot |P|^3)$. Interestingly, it relies on a preprocessing step of time $\mathcal{O}(t \cdot |Q|^3 \cdot |P|^3)$ that is shared among all instances of WIT that we have to check in the second phase. Therefore, the cost of preprocessing only occurs as an additive factor in the overall complexity.

The preprocessing step compiles away the intersection with the memory automaton in $L(S) = L(M) \cap (\text{⧢}_{i\in[1..t]} L(A_i))$. It transforms the thread automata so as to make visible the state changes on the shared memory that their contexts may induce. Formally, we transform each automaton A_i into an *interface automaton* $\text{Int}_M(A_i)$ whose edges are labeled by pairs of states of the memory automaton M. Let $A_i = (P_i, \Sigma \times \{A_i\}, \delta_i, p_i^0, p_i^f)$, we define $\text{Int}_M(A_i)$ as $(P_i, Q \times Q, \Delta_i, p_i^0, p_i^f)$. The new transition relation Δ_i contains all tuples $(p, (q, q'), p')$ where $L(M(q, q')) \cap L(A_i(p, p')) \neq \emptyset$. The interface automata can clearly be constructed in time $\mathcal{O}(t \cdot |Q|^3 \cdot |P|^3)$.

The interface automata do not depend on a memory sequence. To relate them with WIT, we lift $w = q_1 \ldots q_{cs}$ to the alphabet of the interface automata by making explicit the state changes required by the memory sequence: $\text{Int}_M(w) := (q_0, q_1)(q_1, q_2) \ldots (q_{cs}, q_f)$. Finding a word $u \in L(S)$ that induces the memory sequence w now amounts to finding a shuffle of the interface automata that yields $\text{Int}_M(w)$. There is a subtlety. In order to guarantee

cs-many context switches, subsequent letters of $\mathrm{Int}_M(w)$ have to be generated by different interface automata. We use Ш^{cs} to indicate that the shuffle has to switch cs-many times. The proof of the following lemma is straightforward.

▶ **Lemma 4.** *For every $w \in Q^{cs}$, $w \in M(S)$ if and only if $\mathrm{Int}_M(w) \in \text{Ш}_{i\in[1..t]}^{cs} L(\mathrm{Int}_M(A_i))$.*

It remains to decide whether a given word $w \in \Sigma^{cs+1}$ is in the shuffle of given automata $B_1, \ldots, B_t$ with at most $|P|$ states. We present an algorithm for this membership problem that runs in time $\mathcal{O}(4^{cs} \cdot t \cdot cs \cdot |P|)$. This yields the missing factor in the complexity estimation stated in Theorem 2.

▶ **Lemma 5.** *Consider automata $B_1, \ldots B_t$ over Σ with at most $|P|$ states and $w \in \Sigma^{cs+1}$. There is an algorithm to decide $w \in \text{Ш}_{i\in[1..t]}^{cs} L(B_i)$ that runs in time $\mathcal{O}(4^{cs} \cdot t \cdot cs \cdot |P|)$.*

The idea is to construct a graph $G = (V, E)$ with $|V| = \mathcal{O}(2^{cs} \cdot t)$ and $|E| = \mathcal{O}(4^{cs} \cdot t)$ and to reduce the problem to reachability on that graph. We make use of the following observation. The word w is in the shuffle of the B_i if and only if there are disjoint, possibly empty, subwords $w_1, \ldots, w_t$ of w that form a partition of w and that satisfy $w_i \in L(B_i) \cup \{\varepsilon\}$ for all $i \in [1..t]$. The graph G checks this condition by successively covering w by subwords. To this end, it has t stages with $\mathcal{O}(2^{cs})$ vertices each. A vertex in stage i represents a set S_i that contains positions of w that are already covered by subwords. Stage i means that B_i has chosen the last subword. We get an edge between the vertices S_i and S_{i+1} if the subword of w, induced by the positions in $S_{i+1} \setminus S_i$, belongs to $L(B_{i+1})$. We also allow edges not changing the set since an automaton can decide to induce the empty subword. We add an initial vertex s and a final vertex z, where s has edges to the first stage and we only get edges to z if we combined enough subwords so that we cover all positions of w. We prove that we have a path from s to z in G if and only if the word w lies in the shuffle.

Formally, we construct G as follows. The vertices are $V = \{s, z\} \cup \mathcal{P}([1..cs+1]) \times [1..t]$. For the construction of the edges, we assume that an enumeration $\{j_1, \ldots j_m\}$ of a set $S \subseteq [1..cs+1]$ is always in increasing order. Also note that two successive positions of w cannot be covered by one automaton B_i since then we could obtain less than cs context switches. We first add edges from the initial vertex s to $(\emptyset, 1)$ and to all vertices $(S, 1)$, where $S = \{j_1, \ldots, j_m\}$ with $j_l + 1 < j_{l+1}$ for all $l \in [1..(m-1)]$ and such that $w[j_1] \ldots w[j_m] \in L(B_1)$. Note that by $w[j]$, we access the letter of the j-th position of w. For a node (S, i) with $i < t$, we add edges to $(S, i+1)$ and to all vertices $(S', i+1)$, where $S \subseteq S'$ and $S' \setminus S = \{j_1, \ldots, j_m\}$ with $j_l + 1 < j_{l+1}$ for all $l \in [1..(m-1)]$ and such that $w[j_1] \ldots w[j_m] \in L(B_{i+1})$. Finally, we add an edge from $([1..cs+1], t)$ to vertex z. An illustration of the graph is given in Figure 1.

We have $|V| = \mathcal{O}(2^{cs} \cdot t)$ vertices and for each vertex we add at most 2^{cs} edges. Hence, the total number of edges $|E|$ is in $\mathcal{O}(4^{cs} \cdot t)$. Adding an edge takes time $\mathcal{O}(cs \cdot |P|)$ since we have to test membership in one of the B_i. Altogether, we need $\mathcal{O}(4^{cs} \cdot t \cdot cs \cdot |P|)$ time for the construction of G. Checking whether there is a path between s and z can be done in $\mathcal{O}(|V| + |E|) = \mathcal{O}(4^{cs} \cdot t)$ time using a breadth-first search. Hence, we need $\mathcal{O}(4^{cs} \cdot t \cdot cs \cdot |P|)$ time in total. The correctness of the construction follows immediately from the following lemma whose proof can be found in Appendix B.

▶ **Lemma 6.** *Let $S = \{j_1, \ldots, j_m\} \subseteq [1..cs+1]$ and $i \leq t$. There is a path from s to (S, i) if and only if $w[j_1] \ldots w[j_m] \in \text{Ш}_{l\in[1..i]}^{m-1} L(B_l)$.*

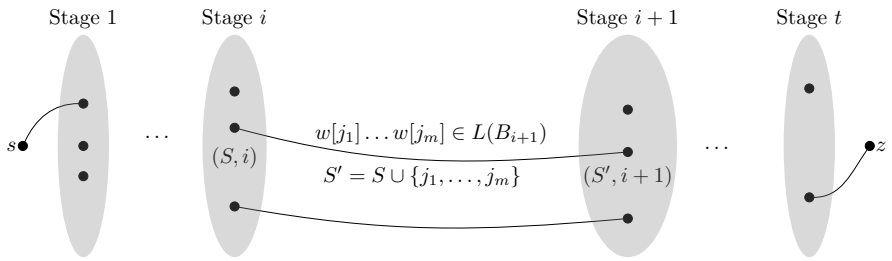

Figure 1 G has a edge between the two vertices (S,i) and $(S',i+1)$ if the subword induced by the positions $S' \setminus S = \{j_1,\dots,j_m\}$ is in the language $L(B_{i+1})$.

3.2 General Models

We extend our algorithm to more general automata models. To this end, we consider classes of automata with polynomial time decidable *membership problem* and polynomial time decidable *regular intersection problem*. The latter problem asks whether $L(M) \cap L(A)$ is non-empty, where M is a finite automaton and A is chosen from the class.

The result that we present in Theorem 7 is particularly interesting for the class of *pushdown automata*. This means that our algorithm can also be used to verify shared memory concurrent programs with recursive threads.

▶ **Theorem 7.** *The problem* $\mathsf{BCS}(cs,|Q|)$ *for threads chosen from a class of automata with polynomial time decidable membership and regular intersection problem, can be solved in* $\mathcal{O}(|Q|^{cs} \cdot 4^{cs} \cdot t \cdot p(n) + t \cdot |Q|^2 \cdot |P|^2 \cdot p(n))$ *time. Here,* $\mathcal{O}(p(n))$ *denotes the runtime of the membership and the regular intersection problem.*

For the proof, we mimic the algorithm presented in Section 3.1 and adjust it where necessary. Let $\mathcal{C}$ be a class as described in the theorem and $(\Sigma, M, (A_i)_{i\in[1..t]})$ an SMCP with threads $A_i \in \mathcal{C}$. First, we construct the interface automata $\mathrm{Int}_M(A_i)$ exactly like in the finite case. Since the construction is based on regular intersection, it can be done for all A_i in $\mathcal{O}(t \cdot |Q|^2 \cdot |P|^2 \cdot p(n))$ time.

The next step is to solve the membership problem for the shuffle language $\coprod^{cs}_{i\in[1..t]} L(B_i)$, where the B_i are from class $\mathcal{C}$. We can adapt the graph construction of Lemma 5. While there is no change in the construction itself, we get a different runtime for adding edges to the graph since this is based on the membership problem of the B_i. Hence, we can construct the graph in $\mathcal{O}(4^{cs} \cdot t \cdot p(n))$ time. Since the correctness is still guaranteed by Lemma 6, this yields the desired complexity.

3.3 Lower bound

We present a lower bound for $\mathsf{BCS}(cs,|Q|)$ of the form $\mathcal{O}^*(2^{cs})$. This shows that the problem is unlikely to be solved in *subexponential time*. Furthermore, we are a $\mathcal{O}(\log|Q|)$ factor in the exponent away from being optimal. The bound is based on ETH. A helpful tool that we make use of is the so-called *Sparsification Lemma* [20] which states that ETH is equivalent to assuming that there is no $2^{o(m)}$-time algorithm that solves m-clause 3-SAT.

▶ **Lemma 8.** *Assuming* ETH, *there is no* $2^{o(cs)}$ *algorithm for* $\mathsf{BCS}(cs,|Q|)$.

Proof. Let $F = C_1 \wedge \dots \wedge C_m$ be a 3-SAT-instance with m clauses and n variables. We show how to construct, in polynomial time, an SMCP $S = (\Sigma, M, (A_i)_{i\in[1..n]})$ with $|Q| = m+1$

so that F has a satisfying assignment if and only if $L_{\leq}(S, m-1) \neq \emptyset$. The latter means that we allow *at most* $m-1$ context switches. Note that an $2^{o(cs)}$-time algorithm for $\mathsf{BCS}(cs, |Q|)$ would induce an $2^{o(cs)}$-time algorithm to determine whether $L_{\leq}(S, cs) \neq \emptyset$. Thus, such an algorithm would induce an $2^{o(m)}$-time algorithm for 3-SAT which contradicts ETH due to the Sparsification Lemma.

We set Σ to be the set of all literals that occur in F. The automaton M has $m+1$ states $q_1, \ldots, q_{m+1}$ with initial state q_1 and final state q_{m+1}. Between the states q_i and q_{i+1}, we get edges labeled by each literal that satisfies clause C_i. We introduce one thread A_i per variable. A_i has three states and its language is $L(A_i) = x_i^+ \cup \neg x_i^+$. Now it is easy to see that F is satisfiable if and only if $L_{\leq}(S, m-1) \neq \emptyset$. ◀

4 Locally bounded memory

In Section 3, we considered the case where the shared memory as a whole was bounded. Now we study a relaxation. For each thread of an SMCP $S = (\Sigma, M, (A_i)_{i\in[1..t]})$, we allocate a particular part of the shared memory M where the thread can toggle certain flags and update its status. We call these allocated parts *local memory states*. Instead of bounding the size of M directly, we will only limit the size of the local memory states. Furthermore, we only allow a thread to hand over control to another thread while M's current state is in the corresponding set of local memory states that belong to the thread.

To formalize the above relaxation, let $S = (\Sigma, M, (A_i)_{i\in[1..t]})$ be a given SMCP with $M = (Q, \Sigma, \delta, q_0, q_f)$, $cs \in \mathbb{N}$ an integer, and $\mu : [1..t] \to \mathcal{P}(Q)$ a map that assigns to each thread the set of local memory states. We call μ the *memory map* and define its *size* to be $|\mu| = \max_{i\in[1..t]} |\mu(i)|$. Furthermore, we define $L(S, cs, \mu)$ to be the language of all words $w \in L(S, cs)$ that have a decomposition $w = w_1 \ldots w_{cs+1}$ so that $w_j \in L(M(q_{j-1}, q_j))$ for all $j \in [1..cs+1]$ and the q_j are states of M that satisfy $q_j \in \mu(\varphi(j))$ for $j \in [1..cs]$ and $q_{cs+1} = q_f$. Now we study the *memory-aware Bounded Context Switching* under-approximation.

Memory-aware Bounded Context Switching (MBCS)	
Input:	An SMCP $S = (\Sigma, M, (A_i)_{i\in[1..t]})$, a bound $cs \in \mathbb{N}$, and a memory map $\mu : [1..t] \to \mathcal{P}(Q)$ (Q denotes the set of states of M).
Question:	Is $L(S, cs, \mu) \neq \emptyset$?

4.1 Intractability

In the case of a bounded memory, we were able to derive an FPT algorithm that did not rely on the number of threads as a parameter. Now that we relax the memory restrictions, we will see that an unbounded number of threads leads to intractability. Formally, we consider the problem $\mathsf{MBCS}(cs, |\mu|)$, which keeps the input of MBCS and introduces the number of context switches and the size of the memory map as parameters.

▶ **Theorem 9.** $\mathsf{MBCS}(cs, |\mu|)$ *is* W[1]*-complete.*

Membership follows immediately from Theorem 1 and the fact $\mathsf{MBCS}(cs, |\mu|) \leq^{fpt} \mathsf{BCS}(cs)$, which is proven in Appendix C. The hardness is obtained by a reduction from the problem of finding a clique of size k in a given graph. This is known to be W[1]-complete [10].

Clique	
Input:	A graph $G = (V, E)$ and an integer $k \in \mathbb{N}$.
Parameter:	k.
Question:	Is there a set of k vertices that forms a clique?

Given a graph $G = (V, E)$ and an integer k, we now show how to reduce *Clique* to $\mathsf{MBCS}(cs, |\mu|)$. We use $\Sigma = \{v, \bar{v}, v^b, v^e \mid v \in V\}$ as the underlying alphabet and construct an SMCP $S = (\Sigma, M, (A_\ell)_{\ell \in V \cup E})$ and a memory map μ so that the number of context switches as well as the size of μ is bounded by an expression only dependent on the parameter k. We call the $(A_v)_{v \in V}$ *vertex automata* and the $(A_e)_{e \in E}$ *edge automata*. The idea behind the reduction is to proceed in two phases. First, we nondeterministically guess k vertices from V, then we verify that the guessed vertices induce a clique by enumerating all edges between the chosen vertices. However, note that the number of such guesses is of size $\mathcal{O}(|V|^k)$. Nevertheless, we can consider each such guess by using the vertex automata A_v as storage.

In the first phase, only automaton M and the vertex automata A_v participate. While M counts up to k, it allows exactly k vertex automata to synchronize with it. Thus, we select the required k vertices. The vertex automata are constructed in such a way that they can synchronize with M at most once in this phase. This ensures that we pick k distinct vertices.

In the second phase, M has to verify that all the selected vertices have edges to each other. This is done by synchronizing with k^2 edge automata and enumerating k^2 edges. However, we have to ensure that the enumerated edges are only edges between the selected vertices. To achieve this, we proceed as follows: for any edge of the form (v, u), M first synchronizes with A_v on the symbol $\bar{v}$ to check if v is among the selected vertices. Then it stores v in its state space and synchronizes with $A_{(v,u)}$ on the letter v^b. This means that the selected edge begins in v. To verify that the edge ends in u, there is a similar synchronization scheme: M and A_u synchronize on $\bar{u}$ and after that M and $A_{(v,u)}$ synchronize on u^e.

The memory map μ allows each vertex automaton A_v to give away control in all states of the first phase and in those states of the second phase, where M stores v in its state space. For an edge automaton $A_{(v,u)}$, the image $\mu(v, u)$ are the remaining states of the second phase. In total, we can guarantee that $|\mu| = \mathcal{O}(k^2)$. A formalization of the construction and the proof of the following lemma can be found in Appendix C.

▶ **Lemma 10.** *Set $cs = k - 1 + 4k^2$. We have that $L(S, cs, \mu) \neq \emptyset$ if and only if G has a clique of size k.*

4.2 A Fixed Parameter Tractability Result

Due to the results of the previous section, we cannot expect to get an FPT algorithm if we do not parametrize by the number of threads. Therefore, we consider the problem $\mathsf{MBCS}(cs, t, |\mu|)$ which introduces as a third parameter the number of threads t. We construct an algorithm, based on the *search tree* technique, to prove the fixed parameter tractability of this parametrization.

▶ **Theorem 11.** $\mathsf{MBCS}(cs, t, |\mu|)$ *is solvable in $\mathcal{O}((t \cdot |\mu|)^{cs+1} \cdot (|Q| \cdot |P|^3 + t \cdot |P|))$ time.*

The idea of the FPT algorithm is to iterate over all possible orders in which the threads may take turns. Having an order fixed, it checks whether there exists a word that respects the order and that witnesses $L(S, cs, \mu) \neq \emptyset$. To this end, we define *interleaving sequences* to be words in $[1..t]^*$ where two subsequent positions are always distinct.

Given an interleaving sequence $\alpha = a_1 \dots a_{cs+1} \in [1..t]^{cs+1}$, we denote by $L_\alpha(S, cs, \mu)$ the language of all $w \in L(S, cs, \mu)$ so that for the decomposition of $w = w_1 \dots w_{cs+1}$ we have that $w_j \in (\Sigma \times \{A_{a_j}\})^+$ or phrased differently: $\varphi(j) = a_j$, for any $j \in [1..cs+1]$. The motivation is that a witness for the non-emptiness of $L(S, cs, \mu)$ clearly obeys some interleaving of the threads.

▶ **Lemma 12.** *$L(S, cs, \mu) \neq \emptyset$ if and only if there exists an interleaving sequence $\alpha \in [1..t]^{cs+1}$ so that $L_\alpha(S, cs, \mu) \neq \emptyset$.*

The equivalence allows us to iterate over all interleaving sequences in $[1..t]^{cs+1}$. Hence, t^{cs+1} is the first factor that appears in the complexity of the algorithm.

Having an *interleaving sequence* fixed, we can focus on solving the following problem. If the sequence α is part of a yes-instance, we call α *valid*.

Validity (VALID)	
Input:	An SMCP $S = (\Sigma, M, (A_i)_{i\in[1..t]})$, a bound $cs \in \mathbb{N}$, a memory map $\mu : [1..t] \to \mathcal{P}(Q)$ and an interleaving sequence $\alpha \in [1..t]^{cs+1}$.
Question:	Is $L_\alpha(S, cs, \mu) \neq \emptyset$?

We construct an algorithm that solves VALID in time $\mathcal{O}(|\mu|^{cs+1} \cdot (|Q| \cdot |P|^3 + t \cdot |P|))$. This yields the total complexity as stated above. Our idea is to build a search tree following the interleaving sequence. In step j, say it is A_{a_j}'s turn, the tree branches along the memory states of A_{a_j} to simulate all partly executions of the thread on M. For each memory state, we compute the set of states of A_{a_j} that synchronizes with the state change on M. At the end, we have to check whether each participating thread can reach a final state.

To be more precise, let $S = (\Sigma, M, (A_i)_{i\in[1..t]})$ be an SMCP with $M = (Q, \Sigma, \delta, q_0, q_f)$ and $A_i = (P_i, \Sigma\times\{A_i\}, \delta_i, p_i^0, p_i^f)$, μ a memory map, and α an interleaving sequence. Without loss of generality, we can assume that $q_f \in \mu(i)$ for any $i \in [1..t]$. We inductively build the following search tree. The root of the tree is labeled by $\mathrm{Init}(\alpha) = (s_1, \dots, s_t, q_0, \alpha)$, where $s_i = \{p_i^0\}$ if i occurs in α and $\perp$ otherwise. Suppose that, in the tree constructed so far, we have a leaf labeled by $(s_1, \dots, s_t, q, a.\beta)$, where $\emptyset \neq s_i \in \mathcal{P}(P_i), q \in Q$, and $a \in [1..t]$. For each $q' \in \mu(a)$ we construct the set $s'_a = \{p' \in P_a \mid \exists p \in s_a \text{ so that } L(M(q, q')) \cap L(A_a(p, p')) \neq \emptyset\}$. If s'_a is non-empty, we add a child node and label it by $(s_1, \dots, s'_a, \dots, s_t, q', \beta)$.

While constructing the tree, the algorithm has to check whether the tree has a leaf labeled by an $\mathcal{F} \in \mathrm{Final}(\alpha) = \{(s_1, \dots s_t, q_f, \varepsilon) \mid s_i \cap \{p_i^f, \perp\} \neq \emptyset \text{ for all } i \in [1..t]\}$. This is due to the fact that checking whether α is valid amounts to checking whether there is a path from the root to a leaf labeled by an element from $\mathrm{Final}(\alpha)$.

▶ **Lemma 13.** *In the tree, there is a path $\mathrm{Init}(\alpha) \to^* \mathcal{F}$, where $\mathcal{F} \in \mathrm{Final}(\alpha)$ if and only if α is valid.*

The resulting tree has outdegree bounded by $|\mu|$ and depth $cs + 1$. Hence, it has at most $|\mu|^{cs+1}$ many nodes in total. For each node, we have to compute the set s'_a. This can be done in $\mathcal{O}(|Q| \cdot |P|^3)$ time. Checking at the end whether the label belongs to $\mathrm{Final}(\alpha)$ can be done in $\mathcal{O}(t \cdot |P|)$. A rough estimation yields the desired complexity of the algorithm.

Proof of Lemma 13. First, let $\alpha = a_1 \dots a_{cs+1}$ be valid. We show that there is a path from $\mathrm{Init}(\alpha)$ to an $\mathcal{F} \in \mathrm{Final}(\alpha)$. From the validity of α, we obtain a $w \in L(S, cs)$ with a decomposition $w = w_1 \dots w_{cs+1}$ and states $q_1, \dots, q_{cs+1}$ of M such that for all $j \in [1..cs+1]$, $q_j \in \mu(a_j)$, $q_{cs+1} = q_f$ and $w_j \in L(M(q_{j-1}, q_j))$. For each $j \in [1..cs+1]$, we show how to construct the sets $s_1^j, \dots, s_t^j$ from which we obtain the required path. For $i \in [1..t]$, we let

$s_i^1 = \{\bot\}$ if i does not occur in α and $s_i^1 = \{p_0^i\}$ otherwise. Now let $j \in [2..cs+1]$ and assume we have already constructed the sets $s_1^{j-1}, \ldots, s_t^{j-1}$. For all $i \in [1..t]$ with $i \neq a_j$, let $s_i^j = s_i^{j-1}$. For $i = a_j$, set $s_i^j = \{p' \mid \exists p \in s_i^{j-1}, L(M(q_{j-1}, q_j)) \cap L(A_i(p, p')) \neq \emptyset\}$. From the existence of w_j, it is easy to see that s_i^j is actually non-empty. Hence, we get that $(s_1^{j-1}, \ldots, s_t^{j-1}, q_{j-1}, a_{j-1} \ldots a_{cs+1}) \to (s_1^j, \ldots, s_t^j, q_j, a_j \ldots a_{cs+1})$ in the tree. Since the projection $w \downarrow \Sigma_i \in L(A_i)$ for all i occurring in α, we get that $p_i^f \in s_i^{cs+1}$ for these i. Hence, we obtain a path $\mathrm{Init}(\alpha) \to^* \mathcal{F}$, where $\mathcal{F} = (s_1^{cs+1}, \ldots, s_t^{cs+1}, q_f, \varepsilon) \in \mathrm{Final}(\alpha)$.

For the converse direction, we assume $\mathrm{Init}(\alpha) \to^* \mathcal{F}$ and let $\alpha = a_1 \ldots a_{cs+1}$. Our goal is to construct non-empty sets $W(j)$ for $j \in [1..cs+1]$ such that the concatenation $W(1) \ldots W(cs+1) \subseteq L_\alpha(S, cs, \mu)$. We start by decomposing the given path into the sequence $\mathrm{Init}(\alpha) = (s_1^1, \ldots, s_t^1, q_0, a_1 \ldots a_{cs+1}) \to \cdots \to (s_1^{cs+1}, \ldots, s_t^{cs+1}, q_f, \varepsilon) = \mathcal{F}$. Now we inductively define sets $P_i(j)$ so that $P_i(j) \cap s_i^j \neq \emptyset$ for any $i \in [1..t]$ and $j \in [1..cs+1]$. Set $P_i(cs+1) = \{p_i^f\}$ if i occurs in α and $P_i(cs+1) = \{\bot\}$ otherwise. Since $\mathcal{F}$ is in $\mathrm{Final}(\alpha)$, we get $P_i(cs+1) \cap s_i^{cs+1} \neq \emptyset$. Suppose, we have already constructed the sets $P_i(j+1)$ with $P_i(j+1) \cap s_i^{j+1} \neq \emptyset$. If $i = a_j$, we define $P_i(j)$ to be the set of states $\{p' \mid \exists p \in P_i(j+1), L(M(q_j, q_{j+1})) \cap L(A_i(p', p)) \neq \emptyset\}$. By construction of the search tree and the fact that $P_i(j+1) \cap s_i^{j+1}$ is non-empty, we also get that $P_i(j) \cap s_i^j$ is non-empty. If $i \neq a_j$, we set $P_i(j) = P_i(j+1)$ and since $s_i^j = s_i^{j+1}$ in this case, we also get $P_i(j) \cap s_i^j \neq \emptyset$. Now we are ready to define the required sets $W(j)$ for a $j \in [1..cs+1]$ as $W(j) = \{w \mid \exists p' \in P_{a_j}(j), p \in P_{a_j}(j+1), w \in L(M(q_j, q_{j+1})) \cap L(A_{a_j}(p', p))\}$. By construction of the $P_i(j)$ and the non-emptiness of the sets, we get that also $W(j)$ is non-empty. Furthermore, it is easy to observe that any word in $W(1) \ldots W(cs+1)$ is a witness for the validity of α. ◀

5 Conclusion

We considered the problem of verifying safety properties on concurrent threads communicating via a shared memory with a bounded context restriction. We reformulated this problem as a non-emptiness of intersection of a regular language with a *bounded* shuffle of languages of arbitrarily many finite state automata.

We showed that bounding the number of threads and the number of context switches is not sufficient to obtain a fixed parameter tractable algorithm. Then, we bounded the size of the shared memory and the number of context switches and derived a FPT algorithm that runs in $\mathcal{O}^*(|Q|^{cs} \cdot 4^{cs})$ time. We extended the algorithm to more general automata models including pushdown systems. Furthermore, based on ETH, we obtained a lower bound of the form $\mathcal{O}^*(2^{cs})$ for the runtime of an algorithm solving the problem.

We also considered a relaxation of bounded memory, where we only bound the number of values each thread can write to the shared memory. In this case, we proved that not parameterizing by the number of participating threads leads to intractability. We then derived an FPT algorithm for the case of a parameterized number of threads that runs in time $\mathcal{O}^*(t^{cs} \cdot |\mu|^{cs})$.

References

1 C. Enea an A. Faran. On atomicity in presence of non-atomic writes. In *TACAS*, 2016.

2 M. F. Atig, A. Bouajjani, K. N. Kumar, and P. Saivasan. On bounded reachability analysis of shared memory systems. In *FSTTCS*, 2014.

3 M. F. Atig, A. Bouajjani, and S. Qadeer. Context-bounded analysis for concurrent programs with dynamic creation of threads. In *TACAS*, 2009.

4 M. F. Atig, A. Bouajjani, and T. Touili. On the reachability analysis of acyclic networks of pushdown systems. In *CONCUR*, 2008.
5 L. Cai, J. Chen, R. G. Downey, and M. R. Fellows. On the parameterized complexity of short computation and factorization. *Arch. Math. Log.*, 1997.
6 M. Cesati. The turing way to parameterized complexity. *J. Comput. Syst. Sci.*, 2003.
7 M. Cygan, J. Nederlof, M. Pilipczuk, M. Pilipczuk, J. M. M. van Rooij, and J. O. Wojtaszczyk. Solving connectivity problems parameterized by treewidth in single exponential time. In *FOCS*, 2011.
8 M. Cygan, M. Pilipczuk, and M. Pilipczuk. Known algorithms for EDGE CLIQUE COVER are probably optimal. In *SODA*, 2013.
9 S. Demri, F. Laroussinie, and P. Schnoebelen. A parametric analysis of the state explosion problem in model checking. In *STACS*, 2002.
10 R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness ii: On completeness for W[1]. *Theoretical Computer Science*, 1995.
11 R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity*. Springer Publishing Company, Incorporated, 2013.
12 A. Durand-Gasselin, J. Esparza, P. Ganty, and R. Majumdar. Model checking parameterized asynchronous shared-memory systems. In *CAV*, 2015.
13 J. Esparza, P. Ganty, and R. Majumdar. Parameterized verification of asynchronous shared-memory systems. In *CAV*, 2013.
14 J. Esparza, P. Ganty, and T. Poch. Pattern-based verification for multithreaded programs. *ACM Trans. Program. Lang. Syst.*, 2014.
15 H. Fernau, P. Heggernes, and Y. Villanger. A multi-parameter analysis of hard problems on deterministic finite automata. *J. Comput. Syst. Sci.*, 2015.
16 H. Fernau and A. Krebs. Problems on finite automata and the exponential time hypothesis. In *CIAA*, 2016.
17 J. Flum and M. Grohe. *Parameterized Complexity Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer-Verlag New York, Inc., 2006.
18 M. Fortin, A. Muscholl, and I. Walukiewicz. On parametrized verification of asynchronous, shared-memory pushdown systems. *CoRR*, 2016.
19 M. Hague. Parameterised pushdown systems with non-atomic writes. In *FSTTCS*, 2011.
20 R. Impagliazzo, R. Paturi, and F. Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, December 2001.
21 D. Kozen. Lower bounds for natural proof systems. In *FOCS*, 1977.
22 S. La Torre, P. Madhusudan, and G. Parlato. Reducing context-bounded concurrent reachability to sequential reachability. In *CAV*, 2009.
23 S. La Torre, P. Madhusudan, and G. Parlato. Model-checking parameterized concurrent programs using linear interfaces. In *CAV*, 2010.
24 A. Lal and T. W. Reps. Reducing concurrent analysis under a context bound to sequential analysis. In *CAV*, 2008.
25 D. Lokshtanov, D. Marx, and S. Saurabh. Slightly superexponential parameterized problems. In *SODA*, 2011.
26 M. Musuvathi and S. Qadeer. Iterative context bounding for systematic testing of multithreaded programs. In *SIGPLAN*, 2007.
27 S. Qadeer and J. Rehof. Context-bounded model checking of concurrent software. In *TACAS*, 2005.
28 T. Wareham. The parameterized complexity of intersection and composition operations on sets of finite-state automata. In *CIAA*, 2000.

A Proofs for Section 2

To prove Theorem 1, we begin by showing that $\mathsf{BCS}(cs,t)$ is a member of $\mathsf{W}[1]$. This is achieved by a parameterized reduction from $\mathsf{BCS}(cs,t)$ to the $\mathsf{W}[1]$-complete problem STMA [5]. After that, we construct a parameterized reduction from $\mathsf{BCS}(cs)$ to $\mathsf{BCS}(cs,t)$ and get that $\mathsf{BCS}(cs)$ is actually in $\mathsf{W}[1]$. For the hardness, we reduce from a parameterized *intersection non-emptiness problem*, which is $\mathsf{W}[1]$-hard.

Short Turing Machine Acceptance (STMA)	
Input:	A nondeterministic Turing machine $\mathcal{TM}$, an input word w, and an integer $k \in \mathbb{N}$.
Parameter:	k.
Question:	Is there a computation of $\mathcal{TM}$ that accepts w in at most k steps?

▶ **Lemma 14.** *We have* $\mathsf{BCS}(cs,t) \leq^{fpt} \mathsf{STMA}$.

Proof. Let $(S = (\Sigma, M, (A_i)_{i\in[1..t]}), cs, t)$ be an instance of $\mathsf{BCS}(cs,t)$ with shared memory $M = (Q, \Sigma, \delta, q_0, q_f)$ and threads $A_i = (P_i, \Sigma \times \{A_i\}, \delta_i, p_i^0, p_i^f)$. We construct a nondeterministic Turing machine $\mathcal{TM}$ and a word w so that $(\mathcal{TM}, w, (cs+1)\cdot t + 2t)$ is an instance of STMA with the property: $L(S, cs) \neq \emptyset$ if and only if there is a computation of $\mathcal{TM}$ accepting w in at most $(cs+1)\cdot t + 2t$ steps.

The idea behind $\mathcal{TM}$ is that the i-th cell of $\mathcal{TM}$'s tape stores the current state of A_i. The states of M and a counter for the turns taken are represented by the control states of $\mathcal{TM}$. Moreover, the transition function of $\mathcal{TM}$ only allows steps which can be carried out simultaneously on M and on one of the A_i. We want $\mathcal{TM}$ to work in three different *modes*: a *switch mode* to perform context switches, a *work mode* to simulate runs of the A_i and M and an *accept mode* which checks if M and those A_i that moved are in a final state.

Formally, we set $\mathcal{TM} = (Q_{\mathcal{TM}}, \Gamma_{\mathcal{TM}}, \delta_{\mathcal{TM}}, q_{init}, q_{acc}, q_{rej})$, where:

- $\Gamma_{\mathcal{TM}} = \dot{\bigcup} P_i \cup \{\mathbf{S}_1, \ldots, \mathbf{S}_t, \mathbf{X}, \$\}$, the $\mathbf{S}_i$ and $\mathbf{X}$ are new letters and \$ is the *left-end marker* of the tape,
- $Q_{\mathcal{TM}} = \{\text{SWITCH}, \text{WORK}\} \times Q \times \{0, \ldots, cs+2\} \cup \{\text{ACCEPT}\} \times \{0, \ldots, t\} \cup \{q_{acc}, q_{rej}\}$ and
- $q_{init} = (\text{SWITCH}, q_0, 0)$.

Moreover, we set $w = \mathbf{S}_1 \ldots \mathbf{S}_t$ and start $\mathcal{TM}$ on this word. Now we will explain the transition function $\delta_{\mathcal{TM}}$. Whenever $\mathcal{TM}$ is in *switch*-mode, a new automaton A_i is chosen to continue (or to start) the run. We allow walking left and right while remembering the current state of M and the number of turns taken but without changing the tape content. So, for $q \in Q$ and $j \leq cs$, we get:

$$((\text{SWITCH}, q, j), p) \to ((\text{SWITCH}, q, j), p, D),$$

where $D \in \{L, R\}$ and $p \in \Gamma_{\mathcal{TM}} \setminus \{\mathbf{X}\}$.
It is also possible to change the mode of $\mathcal{TM}$ to *work*. In this case, we continue the run on the chosen automaton A_i. For $j \leq cs$, we add:

$$((\text{SWITCH}, q, j), p) \to ((\text{WORK}, q, j), p, D).$$

Once $\mathcal{TM}$ is in *work mode*, there are two possibilities. Either the chosen automata A_i did not move before, then there is an $\mathbf{S}_i$ in the currently visited cell, or it has moved before, then the current state of A_i is written in the cell. In the first case, we need a transition rule that *activates* A_i and does a first step. This step has to be synchronized with M. We get

$$((\text{WORK}, q, j), \mathbf{S}_i) \to ((\text{SWITCH}, q', j+1), p', D),$$

for all states $q, q' \in Q$ and $p' \in P_i$ so that $L(M(q,q')) \cap L(A_i(p_i^0, p')) \neq \emptyset$. In the second case, we continue the run on A_i while synchronizing with M:

$$((\text{WORK}, q, j), p) \to ((\text{SWITCH}, q', j+1), p', D),$$

for all states $q, q' \in Q$ and $p, p' \in P_i$ so that $L(M(q,q')) \cap L(A_i(p,p')) \neq \emptyset$.
Note that after changing the mode from *work* to *switch*, we know that a turn was taken and a context switch happened. To track this, $\mathcal{TM}$ increases its counter. This counter is not allowed to go beyond $cs+1$. If this happens, $\mathcal{TM}$ will reject:

$$((\text{SWITCH}, q, cs+2), p) \to q_{rej},$$

for all $q \in Q$ and $p \in \Gamma_{\mathcal{TM}}$.
If $\mathcal{TM}$ arrives in a state of the form $(\text{SWITCH}, q, cs+1)$, where q is a final state of M, then it can enter *accept mode*:

$$((\text{SWITCH}, q, cs+1), p) \to ((\text{ACCEPT}, 0), p, D).$$

Once $\mathcal{TM}$ is in *accept mode*, it moves the head to the left end of the tape via additional moving transitions. Since we assume that the left end is marked by \$, $\mathcal{TM}$ can detect whether it reached the end. We get:

$$((\text{ACCEPT}, 0), p) \to ((\text{ACCEPT}, 0), p, L),$$

for $p \in \Gamma_M \setminus \{\$\}$ and

$$((\text{ACCEPT}, 0), \$) \to ((\text{ACCEPT}, 0), \$, R).$$

After moving to the left end of the tape, $\mathcal{TM}$ will move right and if the current state of A_i, written in the i-th cell, is a final state, it gets replaced by $\mathbf{X}$ and the counter, that counts the number of accepting automata, increases by 1. If $\mathcal{TM}$ sees an $\mathbf{S}_i$ in the i-th cell, it knows that A_i was never activated. This is also counted as accepting. We get:

$$((\text{ACCEPT}, j), p) \to ((\text{ACCEPT}, j+1), \mathbf{X}, D),$$

for p a final state of one of the A_i and

$$((\text{ACCEPT}, j), \mathbf{S}_i) \to ((\text{ACCEPT}, j+1), \mathbf{X}, D).$$

If $\mathcal{TM}$ reads $\mathbf{X}$, it only moves left or right without changing the tape content or counter:

$$((\text{ACCEPT}, j), \mathbf{X}) \to ((\text{ACCEPT}, j), \mathbf{X}, D).$$

When $\mathcal{TM}$ detects t accepting automata, then it will accept:

$$((\text{ACCEPT}, t), \mathbf{X}) \to q_{acc}$$

To simulate $cs+1$ turns of the A_i, $\mathcal{TM}$ needs at most $(cs+1) \cdot t$ steps. Once $\mathcal{TM}$ enters *accept mode*, it needs at most $2t$ steps to verify that each A_i is in a final state or did not move at all. Hence, we are looking for computations of length at most $(cs+1) \cdot t + 2t$. It is easy to observe that the reduction works correctly and can be constructed in polynomial time. ◀

▶ **Lemma 15.** *We have* $\mathsf{BCS}(cs) \leq^{fpt} \mathsf{BCS}(cs, t)$.

Proof. Let $(S = (\Sigma, M, (A_i)_{i\in[1..t]}), cs)$ be an instance of $\mathsf{BCS}(cs)$ with $M = (Q, \Sigma, \delta, q_0, q_f)$ and $A_i = (P_i, \Sigma \times \{A_i\}, \delta_i, p_i^0, p_i^f)$. To construct an instance of $\mathsf{BCS}(cs, t)$, the rough idea is that in cs context switches, we can use at most $cs + 1$ different automata. Hence, we introduce $cs + 1$ new finite automata, where each chooses to simulate one of the A_j.

We set $\Gamma = \Sigma \cup \{\#_1, \ldots, \#_t\} \cup \{\square\}$ and define automaton $B_i = (P_i', \Gamma \times \{B_i'\}, \delta_i', p_i^{ch}, F_i')$ for $i \in [1..cs + 1]$, where:

- $P_i' = \dot{\bigcup}_{j=1}^{t} P_j \cup \{p_i^{ch}, p_i^{lock}\}$ and p_i^{ch}, p_i^{lock} are new states, and
- $F_i' = \dot{\bigcup}_{j=1}^{t} \{p_j^f\} \cup \{p_i^{lock}\}$.

The transition relation δ_i' contains all transition rules of the A_j: if $p \xrightarrow{a} p'$ is an edge in A_j, we get an edge $p \xrightarrow{a} p'$ in B_i. Moreover, we add rules $p_i^{ch} \xrightarrow{\#_j} p_j^0$ for $j \in [1..t]$, and we add $p_i^{ch} \xrightarrow{\square} p_i^{lock}$. An illustration of B_i is given in Figure 2.

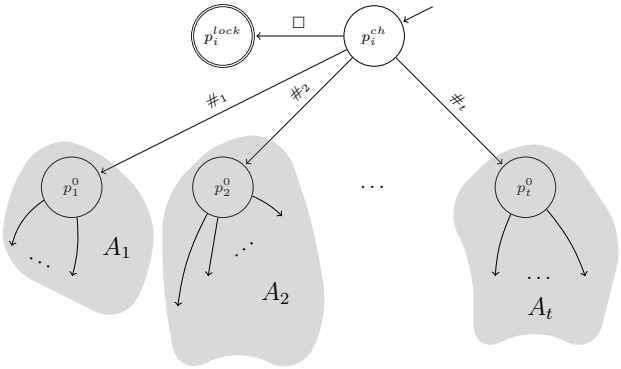

Figure 2 Automaton B_i can either choose to simulate one of the A_j or not to simulate any of the A_j. To this end, it keeps a copy of all A_j and a deadlock state that can be accessed via writing □.

Now we define the finite automaton N to be the tuple $(Q', \Gamma, \delta', q^{ch}, q_f)$, where the states are $Q' = Q \cup \{q_{(\#_j,i)} \mid j \in [1..t] \text{ and } i \in [1..cs]\} \cup \{q^{ch}\}$. The transition relation δ' is the union of the transition relation of M and the rules explained below. To force each B_i to make a choice, we add transitions $q_{(\#_j,i)} \xrightarrow{\#_l} q_{(\#_l,i+1)}$ for $j \in [1..t]$, $j < l \leq t$ and $i \in [1..cs - 1]$. Note that we assume that the choice of the A_l is done in increasing order. This prevents the B_i from choosing the same A_l. For the initial choice, we add transitions $q^{ch} \xrightarrow{\#_l} q_{(\#_l,1)}$ for all $l \in [1..t]$. For the final choice, we add the rules $q_{(\#_j,cs)} \xrightarrow{\#_l} q_0$ for $j \in [1..t]$ and $j < l \leq t$. To give the B_i the opportunity not to simulate any of the A_l, we add the transitions $q_{(\#_j,i)} \xrightarrow{\square} q_{(\#_j,i+1)}$ for $j \in [1..t]$ and $i \in [1..cs - 1]$ and $q_{(\#_j,cs)} \xrightarrow{\square} q_0$ for all $j \in [1..t]$. An image of automaton N can be found in Figure 3.

For the choice of the B_i, we do exactly cs context switches. Then we need another cs context switches for the simulation of the chosen A_l and M. Hence, the tuple given by $(S' = (\Gamma, N, (B_i)_{i\in[1..cs+1]}), 2cs)$ is an instance of $\mathsf{BCS}(cs, t)$ and it is easy to see that we have:

$$L(S', 2cs) \neq \emptyset \text{ if and only if } L(S, cs) \neq \emptyset.$$

◀

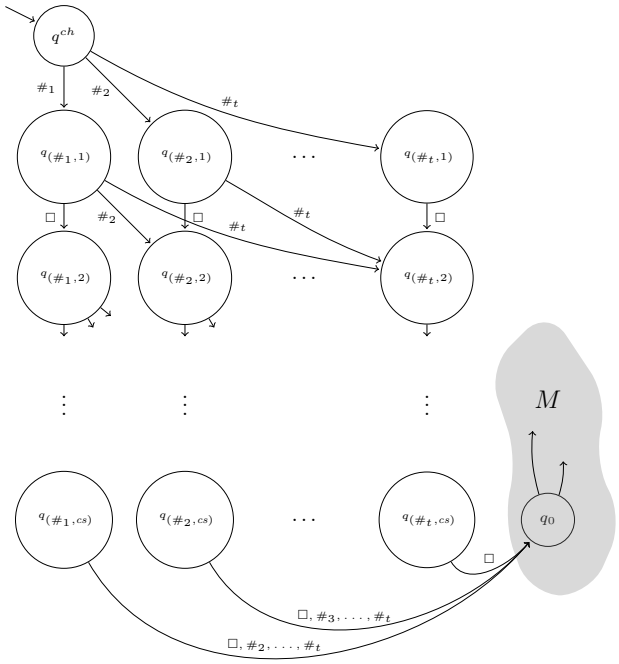

Figure 3 First, automaton N checks whether the B_i have chosen at most $cs{+}1$ different automata A_j to simulate, then it starts the simulation of M.

Bounded DFA Intersection Non-emptiness (BDFAI)	
Input:	Deterministic finite automata $B_1, \ldots, B_n$ over an alphabet Σ and an integer $m \in \mathbb{N}$.
Question:	Does there exist a word w of length m so that $w \in \bigcap_{i=1}^{n} L(B_i)$?

We make use of this bounded version of the *intersection non-emptiness problem*. We parametrize by the length m of the word that we are seeking for and by n, the number of involved automata. We refer to the problem as $\mathsf{BDFAI}(m, n)$ and it is know that $\mathsf{BDFAI}(m, n)$ is $\mathsf{W}[1]$-complete [6, 28].

▶ **Lemma 16.** *We have* $\mathsf{BDFAI}(m, n) \leq^{fpt} \mathsf{BCS}(cs)$.

Proof. Let $(B_1, \ldots, B_n, m)$ be an instance of $\mathsf{BDFAI}(m, n)$ over the alphabet Γ. We construct an instance $(S = (\Sigma, M, (A_i)_{i\in[1..n]}), m \cdot n)$ of $\mathsf{BCS}(cs)$ so that

$$L(S, cs) \neq \emptyset \text{ if and only if } \Sigma^m \cap \bigcap_{i=1}^{n} L(B_i) \neq \emptyset.$$

Set $\Sigma = \Gamma \times \{1, \ldots, n\}$. We construct an automaton A_i which simulates B_i on Γ. To this end, A_i will have the states of B_i and for each transition $p_i \xrightarrow{a} p_i'$ of B_i, we get a transition $p_i \xrightarrow{(a,i)} p_i'$ in A_i. We also set $M = (Q, \Sigma, \delta, q_0, q_f)$, where $Q = \{q_{(a,i)} \mid (a, i) \in \Sigma\} \cup \{q_0\}$, and $q_f = q_0$. For each $a \in \Gamma$, we get transition rules $q_0 \xrightarrow{(a,1)} q_{(a,1)}$, $q_{(a,i)} \xrightarrow{(a,i+1)} q_{(a,i+1)}$ for $i = 1, \ldots, n-2$ and $q_{(a,n-1)} \xrightarrow{(a,n)} q_0$. Automaton M ensures that each B_i reads the same letter and that we only get words in the intersection. Now it is easy to see that the reduction works correctly. ◀

B Proofs for Section 3

Proof of Lemma 6. First, let a path $s \to^+ (S, i)$ be given. We prove this direction by an induction on i. The induction basis $i = 1$ is immediately clear by construction. So, we consider the case $i > 1$. Here, we can split the path into $s \to^+ (S', i-1)$ and $(S', i-1) \to (S, i)$, where $S' = \{a_1, \ldots, a_{m'}\}$. By induction we get that $w[a_1] \ldots w[a_{m'}] \in \text{⧢}_{l \in [1..(i-1)]}^{m'-1} L(B_l)$. Note that we interpret $L(B_1)$ as $\text{⧢}_{l=1}^{m'-1} L(B_l)$ for $i = 2$. Now we distinguish two cases: if $S = S'$, then automaton B_i chooses the empty word and we are done. If $S \neq S'$, then S' is a strict subset of S and we have $S \setminus S' = \{b_1, \ldots, b_{\bar{m}}\}$. By the existence of the edge, we get that $w[b_1] \ldots w[b_{\bar{m}}] \in L(B_i)$ and since the positions $b_1, \ldots, b_{\bar{m}}$ are not successive, we get that $w[j_1] \ldots w[j_m] \in \text{⧢}_{l \in [1..i]}^{m-1} L(B_l)$.

For the other direction, let S be given so that $w[j_1] \ldots w[j_m] \in \text{⧢}_{l \in [1..i]}^{m-1} L(B_l)$. We do an induction on i. The base case $i = 1$ is again clear. So we consider $i > 1$. Let $\{b_1, \ldots, b_{\bar{m}}\}$ be the positions in S that are induced by B_i. This means that we have $w[b_1] \ldots w[b_{\bar{m}}] \in L(B_i)$. We set $S' = S \setminus \{b_1, \ldots, b_{\bar{m}}\}$. If there are no positions of S that are induced by B_i, then we just set $S' = S$. In both cases we get an edge $(S', i-1) \to (S, i)$ and we have $w[a_1] \ldots w[a_{m'}] \in \text{⧢}_{l \in [1..(i-1)]}^{m'-1} L(B_l)$, where $S' = \{a_1, \ldots, a_{m'}\}$. By induction, there is a path $s \to^+ (S', i-1)$ and with the above edge this can be combined to a path $s \to^+ (S, i)$. ◀

C Proofs for Section 4

▶ **Lemma 17.** *We have* $\mathsf{MBCS}(cs, |\mu|) \leq^{fpt} \mathsf{BCS}(cs)$.

Proof. Let $(S = (\Sigma, M, (A_i)_{i \in [1..t]}), cs, \mu)$ be an instance of $\mathsf{MBCS}(cs, \mu)$ with memory $M = (Q, \Sigma, \delta, q_0, q_f)$. First, we set $\Gamma = \Sigma \times \{1, \ldots, t\}$ and multiply the threads A_i by $\{i\}$. This means we keep the structure of A_i and for each transition rule $p \xrightarrow{a} p'$ of A_i, we get a rule $p \xrightarrow{(a,i)} p'$. Denote this automaton by B_i.

Now we construct the new memory N. It keeps t copies of M, one for each thread A_i, and it allows jumps between the copies only out of memory states. Set $N = (\Gamma, Q', \delta', q^{act}, F')$, where

- $Q' = \bigcup_{i=1}^{t} (Q \times \{i\}) \cup \{q^{act}\}$,
- $F' = \bigcup_{i=1}^{t} (\{q_f\} \times \{i\})$,

and δ' consists of the following transition rules:

- The rules of M. For any rule $q \xrightarrow{a} q'$ of M, N has the rules $(q, i) \xrightarrow{(a,i)} (q', i)$ for $i \in [1..t]$.
- Jump transitions. It is possible to switch between the copies of M. For any rule $q \xrightarrow{a} q'$ of M, where $q \in \mu(i)$, we add the rules $(q, i) \xrightarrow{(a,j)} (q', j)$ to N, where $j \neq i$.
- Initial transitions. For a rule $q_0 \xrightarrow{a} q$ of M, we get a rule $q^{act} \xrightarrow{(a,i)} (q, i)$ for $i \in [1..t]$.

We have constructed a BCS-instance $(S' = (\Gamma, N, (B_i)_{i \in [1..t]}), cs)$ and it is easy to observe that

$$L(S', cs) \neq \emptyset \text{ if and only if } L(S, cs, \mu) \neq \emptyset$$

◀

In the following, we present a formalization of the reduction from *Clique* to $\mathsf{MBCS}(cs, |\mu|)$ established in Section 4.1.

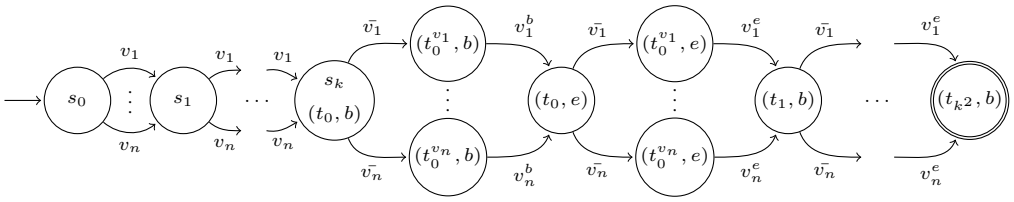

Figure 4 Let $V = \{v_1, \ldots, v_n\}$. Automaton M consists of 2 different phases. In the *guess-phase*, M chooses k vertices. In the *verification-phase*, M enumerates k^2 edges between the vertices. Note that we count an edge (v_i, v_j) twice.

Shared memory. Set $M = (Q, \Sigma, \delta, s_0, q_f)$, where $Q = \text{Phase}_1 \cup \text{Phase}_2$. The states from the set $\text{Phase}_1 = \{s_j \mid j \in [0..k]\}$ are used to select k vertices while the states from the set $\text{Phase}_2 = \{(t_j, b), (t_j^v, b), (t_j, e), (t_j^v, e) \mid j \in [0..k^2], v \in V\}$ are used to enumerate the edges. We also let s_k and (t_0, b) refer to the same node and set $q_f = (t_{k^2}, b)$.

We define $\delta = \delta_1 \cup \delta_2$, where $\delta_1 = \{s_j \xrightarrow{v} s_{j+1} \mid v \in V, j \in [0..(k-1)]\}$ combines the transitions that are used to synchronize with the vertex automata in the first phase. The set δ_2 is the union of $\{(t_j, b) \xrightarrow{\bar{v}} (t_j^v, b), (t_j, e) \xrightarrow{\bar{v}} (t_j^v, e) \mid v \in V, j \in [0..(k^2-1)]\}$, the transitions that ensure that we only enumerate edges among the selected vertices, and $\{(t_j^v, b) \xrightarrow{v^b} (t_j, e), (t_j^v, e) \xrightarrow{v^e} (t_{j+1}, b) \mid v \in V, j \in [0..(k^2-1)]\}$, the transitions for synchronizing with the edge automata. An illustration of M can be found in Figure 4.

Vertex and edge automata. For each $v \in V$, we construct $A_v = (P_v, \Sigma \times \{v\}, \delta_v, p_0^v, p_1^v)$, where $P_v = \{p_0^v, p_1^v\}$ and $\delta_v = \{p_0^v \xrightarrow{v} p_1^v, p_1^v \xrightarrow{\bar{v}} p_1^v\}$. For each $(v, u) \in E$, we build an edge automaton $A_{(v,u)} = (P_{(v,u)}, \Sigma \times \{(v,u)\}, \delta_{(v,u)}, r_0^{(v,u)}, r_2^{(v,u)})$, with set of states $P_{(v,u)} = \{r_0^{(v,u)}, r_1^{(v,u)}, r_2^{(v,u)}\}$ and $\delta_{(v,u)} = \{r_0^{(v,u)} \xrightarrow{v^b} r_1^{(v,u)}, r_1^{(v,u)} \xrightarrow{u^e} r_2^{(v,u)}\}$.

Memory map. For $v \in V$, we define $\mu(v) = \{s_i \mid i \in [1..k]\} \cup \{(t_i^v, b), (t_i^v, e) \mid i \in [0..k^2]\}$. For $(v, u) \in E$, we set $\mu((v, u)) = \{(t_i, b), (t_i, e) \mid i \in [0..k^2]\}$. We get $|\mu| = \mathcal{O}(k^2)$.

Proof of Lemma 10. Let $V = \{v_1, \ldots, v_n\}$ and $E = \{e_1, \ldots, e_m\}$. If $L(S, cs, \mu) \neq \emptyset$ then there is a word w that lies in the shuffle and in the language $L(M)$. Hence w has a decomposition $w = w'.w_1 \ldots w_{k^2}$ where $w' \in (v_1 + \cdots + v_n)^k$ and the words w_i are in the language $(\bar{v_1}.v_1^b + \cdots + \bar{v_n}.v_n^b).(\bar{v_1}.v_1^e + \cdots + \bar{v_n}.v_n^e)$. Note that by construction of the A_v, the word w' consists of k different vertices. Hence, we may assume that $w' = v_1 \ldots v_k$ and we set $V' = \{v_1, \ldots, v_k\}$. We show that the vertices in V' induce a clique. Each w_i is of the form $\bar{v_j}.v_j^b.\bar{v_l}.v_l^e$ with $j, l \leq n$. By using the construction of the A_v again, we get that $j, l \leq k$. Moreover, note that in the part of the shuffle that generates w_i, besides the automata A_{v_j} and A_{v_l}, only one edge automaton can participate: $B_{(v_j, v_l)}$. This means that w_i encodes the edge (v_j, v_l). Since none of the edge automata has loops, they can occur at most once in the whole shuffle. This means that each w_i encodes a different edge between the vertices of V'. Since there are k^2 such edges, we get that V' induces a clique.

For the converse direction, let V' be the set of vertices of a clique of size k. We can assume that $V' = \{v_1, \ldots, v_k\}$. Now the word $w = v_1 \ldots v_k.w_1^1.w_1^2 \ldots w_1^k.w_2^1.w_2^2 \ldots w_k^2 \ldots w_k^k$, where $w_i^j = \bar{v_i}.v_i^b.\bar{v_j}.v_j^e$ is the required witness for $L(S, cs, \mu) \neq \emptyset$. ◀